Compressed Hypothesis Testing: To Mix or Not to Mix?

Myung Cho, Weiyu Xu and Lifeng Lai

arXiv:1609.07528v1 [cs.IT] 23 Sep 2016

Abstract

In this paper, we study the problem of determining k anomalous random variables that have different probability distributions from the rest $(n-k)$ random variables. Instead of sampling each individual random variable separately as in the conventional hypothesis testing, we propose to perform hypothesis testing using mixed observations that are functions of multiple random variables. We characterize the error exponents for correctly identifying the k anomalous random variables under fixed time-invariant mixed observations, random time-varying mixed observations, and deterministic time-varying mixed observations. Our error exponent characterization is through newly introduced notions of *inner conditional Chernoff information* and *outer conditional Chernoff information*. It is demonstrated that mixed observations can strictly improve the error exponents of hypothesis testing, over separate observations of individual random variables. We further characterize the optimal sensing vector maximizing the error exponents, which lead to explicit constructions of the optimal mixed observations in special cases of hypothesis testing for Gaussian random variables. These results show that mixed observations of random variables can reduce the number of required samples in hypothesis testing applications.

Index Terms

compressed sensing, hypothesis testing, Chernoff information, anomaly detection, anomalous random variable, quickest detection

I. Introduction

In many areas of science and engineering such as network tomography, cognitive radio, and radar, one needs to infer statistical information of signals of interest. Statistical information

Myung Cho and Weiyu Xu are with the Department of Electrical and Computer Engineering, The University of Iowa, Iowa City, IA, 52242 USA. Email: {myung-cho, weiyu-xu}@uiowa.edu.

Lifeng Lai is with the Department of Electrical and Computer Engineering, Worcester Polytechnic Institute, Worcester, MA, 01609 USA. Email: llai@wpi.edu.

The work of Lifeng Lai was supported by National Science Foundation through grant CNS-1457076.

of interest can be the mean, variance or even distributions of certain random variables [1–8]. Obtaining such statistical information is essential in detecting anomalous behaviors of random signals. For example, inferring distributions of random variables has important applications in quickest detections of potential hazards, in detecting changes in statistical behaviors of random variables [1, 7, 9], and also in detecting congested links with abnormal delay statistics in network tomography [6].

This paper focuses on the anomaly detection problem. In particular, we consider n random variables, denoted by X_i, $i \in \mathcal{S} = \{1, 2, ..., n\}$, out of which k ($k \ll n$) random variables follow a probability distribution $f_2(\cdot)$ while the much larger set of remaining $(n-k)$ random variables follow another probability distribution $f_1(\cdot)$. However, it is unknown which k random variables follow the distribution $f_2(\cdot)$. Our goal is to infer the subset of random variables that follow $f_2(\cdot)$. In our problem setup, this is equivalent to determining whether X_i follows the probability distribution $f_1(\cdot)$ or $f_2(\cdot)$ for each i. This system model of anomaly detection has appeared in various applications, such as [1, 7, 10–13].

In order to infer the probability distribution of these n random variables, one conventional method is to get l separate samples for each random variable X_i and then use hypothesis testing techniques to determine whether X_i follows the probability distribution $f_1(\cdot)$ or $f_2(\cdot)$ for each i. To ensure correctly identifying the k anomalous random variables with a high probability, at least $\Theta(n)$ samples are needed for hypothesis testing using these samples involving only individual random variables. However, when the number of random variables n grows large, the requirement on sampling rates and sensing resources can be tremendous. For example, in a sensor network, if the fusion center aims to track the anomalies in data generated by n chemical sensors, sending all the data samples of individual sensors to the fusion center will be energy-consuming for the energy-limited sensor network. In such scenarios, we would like to infer the probability distributions of the n random variables with as few data samples as possible. In some applications, due to physical constraints [2, 5, 6], we cannot directly get separate samples of individual random variables. Those difficulties raise the question of whether we can perform hypothesis testing from a much smaller number of samples involving mixed observations.

One way to achieve hypothesis testing with a smaller number of samples is to utilize the sparsity of anomalous random variables, that is the number of anomalous random variables k is much smaller than the total number of random variables n. By utilizing the sparsity of anomalous random variables, [11–14] optimized adaptive *separate* samplings of individual random variables

and reduced the number of needed samples for individual random variables. It is worth noting that the total number of observations is still at least $\Theta(n)$ for these methods, if one is restricted to sample the n random variables individually [11–14].

In this paper, we propose a new approach, named *compressed hypothesis testing*, to find the k anomalous random variables from non-adaptive *mixed* observations of n random variables. In this new approach, instead of a separate observation of an individual random variable, we make each observation a function of the n random variables. Our hypothesis testing method is motivated by compressed sensing [15–17], which is the technique that recovers a *deterministic* sparse vector from its linear projections. Our analysis shows that our new approach can reduce the number of samples required for reliable hypothesis testing. In particular, we show that the number of samples needed to correctly identify the k anomalous random variables can be reduced to $O\left(\frac{k\log(n)}{\min_{p_v,p_w} C(p_v,p_w)}\right)$ observations, where $C(p_v,p_w)$ is the Chernoff information between two possible distributions p_v and p_w for the proposed mixed observations. We also show that mixed observations can strictly increase error exponents of the hypothesis testing, compared with separate sampling of individual random variables. For special cases of Gaussian random variables, we derive optimal mixed measurements to maximize the error exponent of the hypothesis testing. To reduce the computational complexity, we further design efficient algorithms - message passing (MP) based algorithm and Least Absolute Shrinkage and Selection Operator (LASSO) based algorithm - to infer anomalous random variables from mixed observations. We also provide intensive numerical examples to illustrate the advantage of the proposed approach.

This paper is organized as follows. In Section II, we introduce the mathematical models for the considered anomaly detection problem. In Section III-A, we investigate hypothesis testing error performance using time-invariant mixed observations, propose corresponding hypothesis testing algorithms and provide their performance analysis. In Section III-B, we consider using random time-varying mixed observations to identify the anomalous random variables, and derive the error exponent of wrongly identifying the anomalous random variables. In Section III-C, we consider using deterministic time-varying mixed observations for hypothesis testing, and derive a bound on the error probability. In Section IV, we demonstrate, by examples of Gaussian random variable vectors, that linear mixed observations can strictly improve the error exponent over separate sampling of each individual random variables. In Section V, we derive the optimal mixed measurements for Gaussian random variables maximizing the hypothesis testing error exponent. In Section VI, we introduce efficient algorithms to find abnormal random variables using mixed

observations, for large values of n and k. In Section VII, we provide numerical results to demonstrate the effectiveness of our hypothesis testing method from linear measurements. Section VIII describes the conclusion of this paper.

II. Mathematical Models

We consider n independent random variables X_i, $i = 1, \cdots, n$. Out of these n random variables, $(n-k)$ of them follow a known probability distribution $f_1(\cdot)$; while the other k random variables follow another known probability distribution $f_2(\cdot)$:

$$X_i \sim \begin{cases} f_1(\cdot), & i \notin \mathcal{K} \\ f_2(\cdot), & i \in \mathcal{K}, \end{cases} \tag{II.1}$$

where $\mathcal{K} \subset \mathcal{S}$ is an unknown "support" index set, and $|\mathcal{K}| = k \ll n$.

Our goal is to determine $\mathcal{K}$ by identifying those k anomalous random variables with as few samples as possible. We take m mixed observations of the n random variables at m time indices. The measurement at time j is stated as

$$Y^j = g^j(X_1^j, X_2^j, ..., X_n^j),$$

which is a function of n random variables, where $1 \leq j \leq m$. Note that the random variable X_i^j follow the same distribution as the random variable X_i. We assume that the realizations at different time slots are mutually independent. Although our results can be extended to nonlinear observations, in this paper, we specifically consider the case when the functions g^j's are *linear* due to its simplicity and its wide range of applications. When the functions g^j's are linear, the j-th measurement is stated as follows:

$$Y^j = g^j(X_1^j, X_2^j, ..., X_n^j) = \sum_{i=1}^{n} a_i^j X_i^j = \langle a^j, X^j \rangle, \tag{II.2}$$

where a sensing vector $a^j = [a_1^j, a_2^j, \ \dots \ , a_n^j]^T \in \mathbb{R}^{n \times 1}$, and $X^j = [X_1^j, X_2^j, ..., X_n^j]^T$. We obtain an estimate of the index set using a decision function ϕ from Y^j, $j = 1, \cdots, m$, as follows:

$$\hat{\mathcal{K}} = \phi(Y^1, \cdots, Y^m). \tag{II.3}$$

We would like to design the sampling functions g^j's and the decision function ϕ such that the probability

$$P(\hat{\mathcal{K}} \neq \mathcal{K}) \leq \epsilon, \tag{II.4}$$

for an arbitrarily small $\epsilon > 0$.

Our approach is motivated by compressed sensing technique. However, our approach has a major difference from compressed sensing. In our setup, each random variable X_i takes an *independent* realization in each measurement, while in the conventional compressed sensing problem $y = Ax$, where y is the observation vector $y \in \mathbb{R}^{m\times 1}$ and $A \in \mathbb{R}^{m\times n}$ is a sensing matrix, the vector $x = [x_1, x_2, x_3, ..., x_n]^T$ takes the same *deterministic* values across all the measurements. In some sense, our problem is a probabilistic "random-variable" generalization of the compressed sensing problem. In compressed sensing, Bayesian compressed sensing stands out as one model where prior probability distributions of the vector x is considered [18, 19]. However, in [18, 19], the vector $x = [x_1, x_2, x_3, ..., x_n]^T$ is *fixed* once the random variables are realized from the prior probability distribution, and then remains *unchanged* across different measurements. That is fundamentally different from our setting where random variables dramatically change across different measurements. There also exist a collection of research works discussing compressed sensing for *smoothly* time-varying signals [20–22]. In contrast, the objects of interest in this research are random variables taking *completely independent* realizations at different time indices, and we are interested in recovering statistical information of random variables, rather than recovering the deterministic values.

Notations: We denote a random variable and its realization by an uppercase letter and the corresponding lowercase letter respectively. We use X_i to refer to the i-th element of the random variable vector X. We reserve calligraphic uppercase letters $\mathcal{S}$ and $\mathcal{K}$ for index sets, where $\mathcal{S} = \{q : 1 \le q \le n\} = \{q\}_{q=1}^{n}$, and $\mathcal{K} \subseteq \mathcal{S}$. We use superscripts to represent time indices. Hence, x^j represents the realization of a random variable vector X at time j. We reserve the lowercase letters f and p for probability density functions. We also denote the probability density function $p_X(x)$ as $p(x)$ or p_X for simplicity.

III. Compressed Hypothesis Testing

In compressed hypothesis testing, we consider three different types of mixed observations, namely *fixed time-invariant* mixed measurements, *random time-varying* measurements, and *deterministic time-varying* measurements. Table I shows the definition of these types. For these different types of mixed observations, we characterize the number of measurements required to achieve a specified hypothesis testing error probability.

TABLE I

THREE DIFFERENT TYPES OF MIXED OBSERVATIONS

Measurement type	Definition
Fixed time-invariant	The measurement function is the same at every time index.
Random time-varying	The measurement function is randomly generated from a distribution at each time index.
Deterministic time-varying	The measurement function is time-varying but predetermined at each time index.

Algorithm 1: Likelihood ratio test from deterministic time invariant measurements

Data: observation data $y = [y^1, y^2, ..., y^m]^T$, $l = \binom{n}{k}$

Result: k anomalous random variables

1 For each hypothesis H_v ($1 \leq v \leq l$), calculate the likelihood $p_{Y|H_v}(y|H_v)$.

2 Choose the hypothesis with the maximum likelihood.

3 Decide the corresponding k random variables as the anomalous random variables.

A. Fixed Time-Invariant Measurements

In this subsection, we focus on a simple case in which sensing vectors are time-invariant across different time indices, i.e., $a^1 = \cdots = a^m \triangleq a$, $a \in \mathbb{R}^{n\times 1}$. This simple case helps us to illustrate the main idea that will be generalized to more sophisticated schemes in later sections.

We first give the likelihood ratio test algorithm in Algorithm 1, over $l = \binom{n}{k}$ possible hypotheses. To analyze the number of required samples for achieving a certain hypothesis testing error probability, we consider another related hypothesis testing algorithm based on pairwise hypothesis testing in Algorithm 2, which is suboptimal compared to the likelihood ratio test algorithm. There are l possible probability distributions for the output of the function $g(\cdot)$, depending on which k random variables are anomalous. We denote these possible distributions as p_1, p_2, ..., and p_l. Our algorithm is to find the true distribution by doing pairwise Neyman-Pearson hypothesis testing [23] of these l distributions. We provide the complexity for finding the k anomalous random variables by using the time-invariant mixed measurements in Theorem III.1.

Algorithm 2: Hypothesis testing from time-invariant mixed measurements

Data: observation data $y = [y^1, y^2, ..., y^m]^T$

Result: k anomalous random variables

1 For all pairs of distinct probability distributions p_v and p_w ($1 \le v, w \le l$ and $v \ne w$), perform Neyman-Pearson testing for two hypotheses:

- Y^1, Y^2, ..., Y^m follow probability distribution p_v
- Y^1, Y^2, ..., Y^m follow probability distribution p_w

2 Find a certain w^* such that p_{w^*} is the winning probability distribution whenever it is involved in a pairwise hypothesis testing.

3 Declare the k random variables producing p_{w^*} as anomalous random variables.

4 If there is no such w^* in Step 2, then declare a failure in finding the k anomalous random variables.

Theorem III.1. *Consider time-invariant fixed measurements $Y^j = a^{jT} X^j = a^T X^j$, $1 \le j \le m$, for n random variables $X_1, X_2, ..., X_n$. Algorithms 1 and 2 correctly identify the k anomalous random variables with high probability, with $O\left(\frac{k \log(n)}{\min_{1 \le v, w \le l, v \ne w} C(p_v, p_w)}\right)$ mixed measurements. Here, l is the number of hypotheses, p_v, $1 \le v \le l$, is the output probability distribution for measurements $Y = y$ under hypothesis H_v, and*

$$C(p_v, p_w) = -\min_{0 \le \lambda \le 1} \log\left(\int p_v^{\lambda}(y) p_w^{1-\lambda}(y) dy\right)$$

is the Chernoff information between two distributions p_v and p_w.

Proof: In Algorithm 2, for two probability distributions p_v and p_w, we choose the probability likelihood ratio threshold of the Neyman-Pearson testing in such a way that the error probability decreases with the largest possible error exponent, namely the Chernoff information between p_v and p_w:

$$C(p_v, p_w) = -\min_{0 \le \lambda \le 1} \log\left(\int p_v^{\lambda}(y) p_w^{1-\lambda}(y) dy\right).$$

So overall, the smallest possible error exponent of making an error between any pair of probability distributions is

$$E = \min_{1 \le v, w \le l, v \ne w} C(p_v, p_w).$$

Without loss of generality, we assume that p_1 is the true probability distribution for the observation data $Y = y$. Since the error probability P_e in the Neyman-Pearson testing scales

$P_e \triangleq 2^{-mC(p_v,p_w)} \leq 2^{-mE}$, by a union bound over the $l-1$ possible pairs (p_1, p_w), the probability that p_1 is not correctly identified as the true probability distribution scales at most as $l \times 2^{-mE}$, where $l = \binom{n}{k}$. Therefore, $\Theta(k\log(n)E^{-1})$ samplings are enough for identifying the k anomalous samples with high probability.

■

When E grows polynomially with n, this implies a significant reduction in the number of samples needed. If we are allowed to use time-varying non-adaptive sketching functions, we may need fewer samples. In the next subsection, we discuss the performance of time-varying non-adaptive mixed measurements for this problem.

B. Random Time-Varying Measurements

Inspired by compressed sensing where random measurements often give desirable sparse recovery performance [15, 17], we consider random time-varying measurements. In particular, we assume that each measurement is the inner product between X and one independent realization $a^j = [a_1^j, a_2^j, ..., a_n^j]^T$ of a random sensing vector A at time j. Namely, each observation is given by

$$Y^j = \langle A^j, X^j \rangle = \sum_{i=1}^{n} a_i^j X_i^j,\ 1 \leq j \leq m,$$

where $a^j = [a_1^j, a_2^j, ..., a_n^j]^T$ is a realization of the random sensing vector A with pdf $p_A(a)$ at time j. We assume that the realizations a^j's of A are independent across different time indices.

We first give the likelihood ratio test algorithm over $\binom{n}{k}$ hypotheses, namely Algorithm 3. For the purpose of analyzing the error probability of the likelihood ratio test, we further propose a hypothesis testing algorithm based on pairwise comparison, namely Algorithm 4.

The number of samples required to find the abnormal random variables is stated in Theorem III.2.

Theorem III.2. *Consider time-varying random measurements $Y^j = y^j$, $1 \leq j \leq m$, for n random variables X_1, X_2, ..., and X_n. Algorithms 3 and 4 correctly identify the k anomalous random variables with high probability, in $O\left(\frac{k\log(n)}{\min\limits_{1\leq v,w\leq l, v\neq w} IC(p_{A,Y|H_v}, p_{A,Y|H_w})}\right)$ random time-varying measurements. Here, l is the number of hypotheses, $p_{A,Y|H_v}$, $1 \leq v \leq l$, is the output probability distribution for measurements Y and random sensing vectors A under hypothesis H_v, and*

Algorithm 3: Likelihood ratio test from random time-varying measurements

Data: observation and sensing vector at time from 1 to m: $(A^1 = a^1, Y^1 = y^1)$, $(A^2 = a^2, Y^2 = y^2)$, ..., $(A^m = a^m, Y^m = y^m)$

Result: k anomalous random variables

1 For each hypothesis H_v, $1 \le v \le l$, $l = \binom{n}{k}$, calculate the likelihood $p_{A,Y|H_v}(a^j, y^j,\ j = 1, ..., m \mid H_v)$.

2 Choose the hypothesis with the maximum likelihood.

3 Decide the corresponding k random variables as the anomalous random variables.

Algorithm 4: Hypothesis testing from random time-varying measurements

Data: observation and sensing vector at time from 1 to m: $(A^1 = a^1, Y^1 = y^1)$, $(A^2 = a^2, Y^2 = y^2)$, ..., $(A^m = a^m, Y^m = y^m)$

Result: k anomalous random variables

1 For all pairs of hypotheses H_v and H_w ($1 \le v, w \le l$ and $v \ne w$, $l = \binom{n}{k}$), perform Neyman-Pearson testing of the following two hypotheses:

- (A^1, Y^1), (A^2, Y^2), ..., (A^m, Y^m) follow the probability distribution $p_{A,Y|H_v}(a, y|H_v)$
- (A^1, Y^1), (A^2, Y^2), ..., (A^m, Y^m) follow probability distribution $p_{A,Y|H_w}(a, y|H_w)$

2 Find a certain w^* such that H_{w^*} is the winning hypothesis, whenever it is involved in a pairwise hypothesis testing.

3 Declare the k random variables producing H_{w^*} as anomalous random variables.

4 If there is no such w^* in Step 2, declare a failure in finding the k anomalous random variables.

$$
\begin{aligned}
IC(p_{A,Y|H_v}, p_{A,Y|H_w}) &= -\min_{0\le\lambda\le 1} \log\left(\int p_A(a)\left(\int p^{\lambda}_{Y|A,H_v}(y) p^{1-\lambda}_{Y|A,H_w}(y)\, dy\right) da\right) \\
&= -\min_{0\le\lambda\le 1} \log\left(E_A\left(\int p^{\lambda}_{Y|A,H_v}(y) p^{1-\lambda}_{Y|A,H_w}(y)\, dy\right)\right)
\end{aligned}
$$

is the inner conditional Chernoff information *between two hypotheses for measurements Y, conditioned on the probability distribution of time-varying sensing vectors A.*

Proof: In Algorithm 4, for two different hypotheses H_v and H_w, we choose the probability likelihood ratio threshold of the Neyman-Pearson testing in such a way that the hypothesis testing error probability decreases with the largest error exponent, namely the Chernoff information

between $p_{A,Y|H_v}$ and $p_{A,Y|H_w}$:

$$IC(p_{A,Y|H_v}, p_{A,Y|H_w}) = -\min_{0\le\lambda\le1} \log\left(\int p^{\lambda}_{A,Y|H_v}(a,y)p^{1-\lambda}_{A,Y|H_w}(a,y)dy\right).$$

Since the random time-varying sensing vectors are independent of random variable X and the hypothesis H_v or H_w, we obtain the following equations:

$$p_{A,Y|H_v}(a,y|H_v) = p(a|H_v)p(y|H_v,a) = p(a)p(y|H_v,a),$$
$$p_{A,Y|H_w}(a,y|H_w) = p(a|H_w)p(y|H_w,a) = p(a)p(y|H_w,a).$$

Then the Chernoff information is simplified to

$$\begin{aligned} IC(p_{A,Y|H_v}, p_{A,Y|H_w}) &= -\min_{0\le\lambda\le1} \log\left(\int p_A(a)\left(\int p^{\lambda}_{Y|A,H_v}(y)p^{1-\lambda}_{Y|A,H_w}(y)\,dy\right)da\right) \\ &= -\min_{0\le\lambda\le1} \log\left(E_A\left(\int p^{\lambda}_{Y|A,H_v}(y)p^{1-\lambda}_{Y|A,H_w}(y)\,dy\right)\right) \end{aligned}$$

Using Holder's inequality, we have

$$IC(p_{A,Y|H_v}, p_{A,Y|H_w}) \ge -\min_{a} \log\left(1 - p_A(a) + p_A(a)e^{-C(p_{Y|A,H_v}, p_{Y|A,H_w})}\right),$$

where

$$C(p_{Y|A,H_v}, p_{Y|A,H_w}) = -\min_{0\le\lambda\le1} \log\left(\int p^{\lambda}_{Y|A,H_v}(y)p^{1-\lambda}_{Y|A,H_w}(y)dy\right)$$

is the well-known Chernoff information between $p_{Y|A,H_v}$, and $p_{Y|A,H_j}$. So as long as there exist sensing vectors A of a positive probability, such that the regular Chernoff information is positive, then the inner condition Chernoff information $IC(p_{A,Y|H_v}, p_{A,Y|H_w})$ will also be positive.

Overall, the smallest possible error exponent between any pair of hypotheses is

$$E = \min_{1\le v,w\le l, v\ne w} IC(p_{A,Y|H_v}, p_{A,Y|H_w}).$$

Without loss of generality, we assume H_1 is the true hypothesis. Since the error probability P_e in the Neyman-Pearson testing is

$$P_e \triangleq 2^{-m(IC(p_{A,Y|H_v}, p_{A,Y|H_w}))} \le 2^{-mE}.$$

By a union bound over the $l-1$ possible pairs (H_1, H_w), the probability that H_1 is not correctly identified as the true hypothesis is upper bounded by $l \times 2^{-mE}$ in terms of scaling. So $m = \Theta(k\log(n)E^{-1})$ samplings are enough for identifying the k anomalous samples with high probability. When E grows polynomially with n, this implies a significant reduction in the number of required samples. ■

Algorithm 5: Likelihood ratio test from deterministic time-varying measurements

Data: observation data $y = [y^1, y^2, ..., y^m]^T$, $y \in \mathbb{R}^{m\times 1}$, deterministic sensing vectors $[a^1, a^2, ..., a^m]^T \in \mathbb{R}^{m\times n}$

Result: k anomalous random variables

1 For each hypothesis H_v $(1 \le v \le l)$, calculate the likelihood $p_{Y|H_v,A}(y|H_v, a)$.

2 Choose the hypothesis with the maximum likelihood.

3 Declare the corresponding k random variables as the anomalous random variables.

C. Deterministic Time-Varying Measurements

In this subsection, we consider mixed measurements which are allowed to vary over time. However, each sensing vector is predetermined, so that exactly $P(A = a)m$ (assuming that $P(A = a)m$ are integers) measurements use a realized sensing vector a. In contrast, in random time-varying measurements, each sensing vector A is taken randomly, and thus the number of measurements taking realization a is random. We define the predetermined sensing vector at time j as a^j.

For deterministic time-varying measurements, we first give the likelihood ratio test algorithm among $l = \binom{n}{k}$ hypotheses (Algorithm 5). To analyze the error probability, we consider another hypothesis testing algorithm based on pairwise comparison with deterministic time-varying measurements (Algorithm 6).

Theorem III.3. *Consider time-varying deterministic observations $Y^j = y^j$, $1 \le j \le m$, for n random variables X_1, X_2, ..., and X_n. l is the number of hypotheses for the distribution of the vector $X = [X_1, X_2, ..., X_n]^T$.*

For $\lambda \in [0, 1]$ and two hypotheses H_v and H_w $(1 \le v, w \le l)$, define

$$p_\lambda(y|a, H_v, H_w) = \frac{p^\lambda_{Y|A,H_v}(y|a, H_v)p^{1-\lambda}_{Y|A,H_w}(y|a, H_w)}{\int p^\lambda_{Y|A,H_v}(y|a, H_v)p^{1-\lambda}_{Y|A,H_w}(y|a, H_w)dy},$$

$$Q_{\lambda,v\to w} = E_A\left\{D\left(p_\lambda(y|a, H_v, H_w) \parallel p(y|a, H_v)\right)\right\},$$

$$Q_{\lambda,w\to v} = E_A\left\{D\left(p_\lambda(y|a, H_v, H_w) \parallel p(y|a, H_w)\right)\right\}.$$

Furthermore, we define the outer conditional Chernoff information $OC(p_{Y|A,H_v}, p_{Y|A,H_w})$ *between H_v and H_w, under deterministic time-varying sensing vector A, as*

$$OC(p_{Y|A,H_v}, p_{Y|A,H_w}) = Q_{\lambda,v\to w} = Q_{\lambda,w\to v},$$

Algorithm 6: Hypothesis testing from deterministic time-varying measurements

Data: observation data $y = [y^1, y^2, ..., y^m]^T$, $y \in \mathbb{R}^{m\times 1}$, deterministic sensing vectors $[a^1, a^2, ..., a^m]^T \in \mathbb{R}^{m\times n}$

Result: k anomalous random variables

1 For all pairs of hypotheses H_v and H_w ($1 \le v, w \le l$ and $v \ne w$), perform Neyman-Pearson testing of the following two hypotheses:

- Y^1, Y^2, ..., Y^m follow the probability distribution $p_{Y|H_v,A}(y|H_v, a)$
- Y^1, Y^2, ..., Y^m follow probability distribution $p_{Y|H_w,A}(y|H_w, a)$

2 Find a certain w^*, such that H_{w^*} is the winning hypothesis, whenever it is involved in a pairwise hypothesis testing.

3 Declare the k random variables producing H_{w^*} as anomalous random variables.

4 If there is no such w^* in Step 2, declare a failure in finding the k anomalous random variables.

where λ is chosen such that $Q_{\lambda,v\to w} = Q_{\lambda,w\to v}$.

Then with $O(\frac{k\log(n)}{\min\limits_{1\le v,w\le l, v\ne w} OC(p_{Y|A,H_v}, p_{Y|A,H_w})})$ random time-varying measurements, with high probability, Algorithms 5 and 6 correctly identify the k anomalous random variables. Here l is the number of hypotheses, $p_{Y|A,H_v}$, $1 \le v \le l$ is the output probability distribution for observations Y under hypothesis H_v and sensing vector A, and $OC(p_{Y|A,H_v}, p_{Y|A,H_w})$ is the outer conditional Chernoff information.

Moreover, the outer conditional Chernoff information is equal to

$$\begin{aligned} OC(p_{Y|A,H_v}, p_{Y|A,H_w}) &= -\min_{0\le\lambda\le 1} \int p_A(a) \log\left(\int p^{\lambda}_{Y|A,H_v}(y) p^{1-\lambda}_{Y|A,H_w}(y)\, dy\right) da \\ &= -\min_{0\le\lambda\le 1} E_A\left(\log\left(\int p^{\lambda}_{Y|A,H_v}(y) p^{1-\lambda}_{Y|A,H_w}(y)\, dy\right)\right). \end{aligned}$$

We place the proof of this theorem in the appendix.

IV. Examples of Compressed Hypothesis Testing

In this section, we provide simple examples in which smaller error probability can be achieved in hypothesis testing through mixed observations than the traditional individual sampling approach,

under the same number of measurements. In particular, we consider Gaussian distributions in our examples.

A. Example 1: two Gaussian random variables

In this example, we consider $n = 2$, and $k = 1$. We group the two random variables X_1 and X_2 in a random vector $[X_1, X_2]^T$. Suppose that there are two hypotheses for a 2-dimensional random vector $[X_1, X_2]^T$, where X_1 and X_2 are independent:

- H_1: $X_1 \sim \mathcal{N}(A, \sigma^2)$ and $X_2 \sim \mathcal{N}(B, \sigma^2)$,
- H_2: $X_1 \sim \mathcal{N}(B, \sigma^2)$ and $X_2 \sim \mathcal{N}(A, \sigma^2)$.

Here A and B are two distinct constants, and σ^2 is the variance of the two Gaussian random variables. At each time index, only one observation is allowed, and the observation is restricted to a linear mixing of X_1 and X_2. Namely

$$Y^j = a_1 X_1 + a_2 X_2.$$

We assume that the sensing vector $[a_1, a_2]^T$ does not change over time.

Clearly, when $a_1 \neq 0$ and $a_2 = 0$, the sensing vector reduces to a separate observation of X_1; and when $a_1 = 0$ and $a_2 \neq 0$, it reduces to a separate observation of X_2. In both these two cases, the observation follows distribution $\mathcal{N}(A, \sigma^2)$ for one hypothesis, and follows distribution $\mathcal{N}(B, \sigma^2)$ for the other hypothesis. The Chernoff information between these two distributions are

$$C(\mathcal{N}(A, \sigma^2), \mathcal{N}(B, \sigma^2)) = \frac{(A-B)^2}{8\sigma^2}.$$

When hypothesis H_1 holds, the observation Y^j follows Gaussian distribution $\mathcal{N}(a_1 A + a_2 B, (a_1^2 + a_2^2)\sigma^2)$. Similarly, when hypothesis H_2 holds, the observation Y^j follows Gaussian distribution $\mathcal{N}(a_1 B + a_2 A, (a_1^2 + a_2^2)\sigma^2)$. The Chernoff information between these two Gaussian distributions $\mathcal{N}(a_1 A + a_2 B, (a_1^2 + a_2^2)\sigma^2)$ and $\mathcal{N}(a_1 B + a_2 A, (a_1^2 + a_2^2)\sigma^2)$ is given by

$$\frac{[(a_1 A + a_2 B) - (a_1 B + a_2 A)]^2}{8(a_1^2 + a_2^2)\sigma^2} = \frac{[(a_1 - a_2)^2 (A-B)^2]}{8(a_1^2 + a_2^2)\sigma^2} \leq \frac{2(A-B)^2}{8\sigma^2},$$

where the last inequality follows from Cauchy-Schwarz inequality, and takes equality when $a_1 = -a_2$.

Compared with the Chernoff information for separate observations of X_1 or X_2, the linear mixing of X_1 and X_2 doubles the Chernoff information. This shows that linear mixed observations

can offer strict improvement in terms of reducing the error probability in hypothesis testing, by increasing the error exponent.

B. Example 2: Gaussian random variables with different means

In this example, we consider the mixed observations for two hypotheses of Gaussian random vectors. In general, suppose that there are two hypotheses for an n-dimensional random vector $[X_1, X_2, ..., X_n]^T$,

- H_1: $[X_1, X_2, ..., X_n]$ follow jointly Gaussian distributions $\mathcal{N}(\mu_1, \Sigma_1)$,
- H_2: $[X_1, X_2, ..., X_n]$ follow jointly Gaussian distributions $\mathcal{N}(\mu_2, \Sigma_2)$.

Here Σ_1 and Σ_2 are both $n \times n$ covariance matrices.

At each time instant, only one observation is allowed, and the observation is restricted to a time-invariant sensing vector. Namely

$$Y^j = \langle A, X^j \rangle,$$

where $A \in \mathbb{R}^{n\times 1}$.

Under these conditions, the observation follows distribution $\mathcal{N}(A^T\mu_1, A^T\Sigma_1 A)$ for hypothesis H_1, and follows distribution $\mathcal{N}(A^T\mu_2, A^T\Sigma_2 A)$ for the other hypothesis H_2. We would like to choose a sensing vector A which maximizes the Chernoff information between the two possible univariate Gaussian distributions, namely

$$\max_A C(\mathcal{N}(A^T\mu_1, A^T\Sigma_1 A), \mathcal{N}(A^T\mu_2, A^T\Sigma_2 A)).$$

In fact, the Chernoff information between these two distributions [24] is

$$\begin{aligned} &C(\mathcal{N}(A^T\mu_1, A^T\Sigma_1 A), \mathcal{N}(A^T\mu_2, A^T\Sigma_2 A)) \\ &= \max_{0\le\alpha\le 1} \left[\frac{1}{2} \log \left(\frac{A^T(\alpha\Sigma_1 + (1-\alpha)\Sigma_2)A}{(A^T\Sigma_1 A)^\alpha (A^T\Sigma_2 A)^{1-\alpha}} \right) + \frac{\alpha(1-\alpha)(A^T(\mu_1-\mu_2))^2}{2A^T(\alpha\Sigma_1 + (1-\alpha)\Sigma_2)A} \right]. \end{aligned}$$

We first look at the special case when $\Sigma = \Sigma_1 = \Sigma_2$. Under this condition, the maximum Chernoff information is given by

$$\max_A \max_{0\le\alpha\le 1} \frac{\alpha(1-\alpha)[A^T(\mu_1-\mu_2)]^2}{2A^T\Sigma A}.$$

Taking $A' = \Sigma^{\frac{1}{2}} A$, then this reduces to

$$\max_{A'} \max_{0\le\alpha\le 1} \frac{\alpha(1-\alpha)[(A')^T\Sigma^{-\frac{1}{2}}(\mu_1-\mu_2)]^2}{2(A')^T A'}.$$

From Cauchy-Schwarz inequality, it is easy to see that the optimal $\alpha = \frac{1}{2}$, $A' = \Sigma^{-\frac{1}{2}}(\mu_1 - \mu_2)$, and $A = \Sigma^{-1}(\mu_1 - \mu_2)$. Under these conditions, the maximum Chernoff information is given by

$$\frac{1}{8}(\mu_1 - \mu_2)^T \Sigma^{-1}(\mu_1 - \mu_2).$$

Note that in general, $A' = \Sigma^{-\frac{1}{2}}(\mu_1 - \mu_2)$ is not a separate observation of a certain individual random variable, but rather a linear mixing of the n random variables.

C. Example 3: Gaussian random variables with different variances

In this example, we look at the mixed observations for Gaussian random variables with different variances. Consider the same setting in Example 2, except that we now look at the special case when $\mu = \mu_1 = \mu_2$. We will study the optimal sensing vector under this scenario. Then the Chernoff information becomes

$$C(\mathcal{N}(A^T\mu, A^T\Sigma_1 A), \mathcal{N}(A^T\mu, A^T\Sigma_2 A)) = \max_{0\le\alpha\le 1} \frac{1}{2}\log\left(\frac{A^T(\alpha\Sigma_1 + (1-\alpha)\Sigma_2)A}{(A^T\Sigma_1 A)^\alpha (A^T\Sigma_2 A)^{1-\alpha}}\right).$$

To find the optimal sensing vector A, we are solving this optimization problem

$$\max_A \max_{0\le\alpha\le 1} \frac{1}{2}\log\left(\frac{A^T(\alpha\Sigma_1 + (1-\alpha)\Sigma_2)A}{(A^T\Sigma_1 A)^\alpha (A^T\Sigma_2 A)^{1-\alpha}}\right).$$

For a certain A, we define

$$B = \frac{\max\left(A^T\Sigma_1 A, A^T\Sigma_2 A\right)}{\min\left(A^T\Sigma_1 A, A^T\Sigma_2 A\right)}.$$

Note that $B \ge 1$. By symmetry over α and $1-\alpha$, maximizing the Chernoff information can always be reduced to

$$\max_{B\ge 1} \max_{0\le\alpha\le 1} \frac{1}{2}\log\left(\frac{\alpha + (1-\alpha)B}{B^{1-\alpha}}\right). \tag{IV.1}$$

The optimal α is obtained as follows by finding the point which makes the first order differential equation to zero:

$$\alpha = \frac{-(B-1) + B\log(B)}{(B-1)\log(B)}.$$

By plugging the obtained optimal α to (IV.1), we obtain the following optimization problem:

$$\max_{B\ge 1} \frac{1}{2}\left\{-1 + \frac{B}{B-1}\log(B) + \log\left(\frac{B-1}{B\log(B)}\right)\right\}. \tag{IV.2}$$

We note that the objective function is an increasing function in B, when $B \ge 1$.

Lemma IV.1. *The optimal objective value of the following optimization problem*

$$\max_{0\le\alpha\le 1} \frac{1}{2}\log\left(\frac{\alpha+(1-\alpha)B}{B^{1-\alpha}}\right),$$

is an increasing function in $B \ge 1$.

Proof: We only need to show that for any $\alpha \in [0,1]$, $\left(\frac{\alpha+(1-\alpha)B}{B^{1-\alpha}}\right)$ is an increasing function in $B \ge 1$. In fact, the derivative of it with respect to α is

$$\alpha(1-\alpha)(B^{\alpha-1}-B^{\alpha-2}) \ge 0.$$

Then the conclusion of this lemma immediately follows. ■

This means we need to maximize B, in order to maximize the Chernoff information. Hence, to find the optimal A maximizing B, we solve the following two optimization problems:

$$\max_{A} A^T\Sigma_1 A \qquad \text{subject to} \quad A^T\Sigma_2 A \le 1, \tag{IV.3}$$

and

$$\max_{A} A^T\Sigma_2 A \qquad \text{subject to} \quad A^T\Sigma_1 A \le 1. \tag{IV.4}$$

Then the maximum of the two optimal objective values is equal to the optimal objective value of optimizing B, and the corresponding A is the optimal sensing vector maximizing the Chernoff information. These two optimization problems are not convex optimization programs, however, they still admit zero duality gap from the S-procedure, and can be efficiently solved [25]. In fact, they are respectively equivalent to the following two semidefinite programming optimization problems:

$$\begin{aligned} &\underset{\gamma,\lambda}{\text{minimize}} \quad -\gamma \\ &\text{subject to} \quad \lambda \ge 0 \\ &\qquad\qquad \begin{pmatrix} -\Sigma_1+\lambda\Sigma_2 & 0 \\ 0 & -\lambda-\gamma \end{pmatrix} \ge 0, \end{aligned} \tag{IV.5}$$

and

$$\begin{aligned} &\underset{\gamma,\lambda}{\text{minimize}} \quad -\gamma \\ &\text{subject to} \quad \lambda \ge 0 \\ &\qquad\qquad \begin{pmatrix} -\Sigma_2+\lambda\Sigma_1 & 0 \\ 0 & -\lambda-\gamma \end{pmatrix} \ge 0. \end{aligned} \tag{IV.6}$$

Thus they can be efficiently solved.

For example, when Σ_1 and Σ_2 are given as follows:

$$\Sigma_1 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 100 \end{bmatrix}, \qquad \Sigma_2 = \begin{bmatrix} 100 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix},$$

we obtain $[0, 0, -1]^T$ for the optimal sensing vector A. This is because the diagonal elements in the covariance matrices Σ_1 and Σ_2 represent the variance of random variables $[X_1, X_2, X_3]^T$, and the biggest difference in variance is shown in random variable X_1 or X_3. Therefore, checking the random variable X_3 (or X_1) through realizations is an optimal way to figure out whether the random variables follows Σ_1 or Σ_2. On the other hand, if the random variables are dependent, and Σ_1 and Σ_2 are given as follows:

$$\Sigma_1 = \begin{bmatrix} 1 & 0.5 & 0.5 \\ 0.5 & 1 & 0.5 \\ 0.5 & 0.5 & 100 \end{bmatrix}, \qquad \Sigma_2 = \begin{bmatrix} 100 & 0.5 & 0.5 \\ 0.5 & 1 & 0.5 \\ 0.5 & 0.5 & 1 \end{bmatrix},$$

then, we obtain $[0.4463, 0.0022, -0.8949]^T$ for the optimal sensing vector A which maximizes the Chernoff information.

D. Example 4: $k = 1$ anomalous random variable among $n = 7$ random variables

Consider another example, where $k = 1$ and $n = 7$. $n - k = 6$ random variables follow the distribution $\mathcal{N}(0, 1)$; and the other random variable follows distribution $\mathcal{N}(0, \sigma^2)$, where $\sigma^2 \neq 1$. We assume that all random variables X_1, X_2,..., and X_7 are independent. So overall, there are 7 hypotheses:

- H_1: (X_1, X_2, ..., X_7) ~ ($\mathcal{N}(0, \sigma^2)$, $\mathcal{N}(0, 1)$, ..., $\mathcal{N}(0, 1)$),
- H_2: (X_1, X_2, ..., X_7) ~ ($\mathcal{N}(0, 1)$, $\mathcal{N}(0, \sigma^2)$, ..., $\mathcal{N}(0, 1)$),
- .....
- H_7: (X_1, X_2, ..., X_7) ~ ($\mathcal{N}(0, 1)$, $\mathcal{N}(0, 1)$, ..., $\mathcal{N}(0, \sigma^2)$).

We first assume that separate observations of these 7 random variables are obtained. For any pair of hypotheses H_v and H_w, the probability distributions for the output are respectively $\mathcal{N}(0, \sigma^2)$ and $\mathcal{N}(0, 1)$, when X_v is observed; the probability distributions for the output are

respectively $\mathcal{N}(0,1)$ and $\mathcal{N}(0,\sigma^2)$, when X_w is observed. From (VIII.1), the Chernoff information between H_v and H_w is given by

$$\begin{aligned} OC(p_{Y|A,H_v}, p_{Y|A,H_w}) &= -\min_{0\le\lambda\le1} E_A\left(\log\left(\int p^{\lambda}_{Y|A,H_v}(y)p^{1-\lambda}_{Y|A,H_v}(y)\,dy\right)\right) \\ &= -\min_{0\le\lambda\le1}\left[\frac{1}{7}\left(\log\left(\int p^{\lambda}_{\mathcal{N}(0,1)}(y)p^{1-\lambda}_{\mathcal{N}(0,\sigma^2)}(y)\,dy\right)\right)\right. \\ &\quad \left.+\frac{1}{7}\left(\log\left(\int p^{\lambda}_{\mathcal{N}(0,\sigma^2)}(y)p^{1-\lambda}_{\mathcal{N}(0,1)}(y)\,dy\right)\right)\right]. \end{aligned}$$

Optimizing over λ, we obtain the optimal $\lambda = \frac{1}{2}$, and that

$$OC(p_{Y|A,H_v}, p_{Y|A,H_w}) = \frac{1}{7}\log(\frac{B+1}{2B^{\frac{1}{2}}}), \tag{IV.7}$$

where $B = \frac{\max(\sigma^2,1)}{\min(\sigma^2,1)}$.

Now we consider using the parity check matrix of $(7,4)$ Hamming codes to obtain measurements as follows:

$$\begin{bmatrix} a^{1T} \\ a^{2T} \\ a^{3T} \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 1 & 1 & 0 & 1 \\ 0 & 1 & 0 & 1 & 0 & 1 & 1 \\ 0 & 0 & 1 & 0 & 1 & 1 & 1 \end{bmatrix}.$$

We use the row vectors of the parity check matrix as sensing vectors for deterministic time-varying measurements. For example, the i-th row vector is used for the j-th mixed measurement, where $i = (j \bmod 3)+1$. For any pair of hypotheses H_v and H_w, there is always a sensing vector among 3 sensing vectors which measures one and only one of X_v and X_w. Without loss of generality, we assume that that a sensing vector measures X_v but not X_w. Suppose H_v is true, then the output probability distribution for that measurement is $\mathcal{N}(0,\sigma^2+s(A))$; otherwise when H_w is true, the output probability distribution is given by $\mathcal{N}(0,1+s(A))$, where $1+s(A)$ is the number of ones in the sensing vector A. The Chernoff information between H_v and H_w is bounded by

$$\begin{aligned} OC(p_{Y|A,H_v}, p_{Y|A,H_w}) &= -\min_{0\le\lambda\le1} E_A\left(\log\left(\int p^{\lambda}_{Y|A,H_v}(y)p^{1-\lambda}_{Y|A,H_v}(y)\,dy\right)\right) \\ &\ge -\min_{0\le\lambda\le1}\frac{1}{3}\log\left(\int p^{\lambda}_{\mathcal{N}(0,\sigma^2+s(A))}(y)p^{1-\lambda}_{\mathcal{N}(0,1+s(A))}(y)\,dy\right), \end{aligned}$$

where the final inequality is obtained by considering only one sensing vector among 3 sensing vectors. From (IV.2), this lower bound is given by

$$\frac{1}{6}\left\{-1+\frac{B}{B-1}\log(B)+\log\left(\frac{B-1}{B\log(B)}\right)\right\},$$

where

$$B = \frac{\max(\sigma^2+s(A),1+s(A))}{\min(\sigma^2+s(A),1+s(A))}.$$

Simply taking $\lambda = \frac{1}{2}$, we obtain another lower bound of $OC(p_{Y|A,H_v}, p_{Y|A,H_w})$ as follows:

$$\begin{aligned} & -\min_{0\le\lambda\le 1} \frac{1}{3}\log\left(\int p^{\lambda}_{\mathcal{N}(0,\sigma^2+s(A))}(y) p^{1-\lambda}_{\mathcal{N}(0,1+s(A))}(y)\, dy\right) \\ & \ge \begin{cases} \frac{1}{3}\times\frac{1}{2}\log\left(\frac{\sigma^2+s(A)+1}{2\sqrt{\sigma+s(A)}}\right), & \text{if } \sigma^2 \ge 1, \\ \frac{1}{3}\times\frac{1}{2}\log\left(\frac{s(A)+2}{2\sqrt{1+s(A)}}\right), & \text{if } \sigma^2 < 1. \end{cases} \end{aligned}$$

When $\sigma^2 \gg 1$, for separate observations,

$$OC(p_{Y|A,H_v}, p_{Y|A,H_w}) = \frac{1}{7}\log(\frac{\sigma^2+1}{2\sigma}) \sim \frac{1}{7}\log(\frac{\sigma}{2});$$

while for measurements through the parity-check matrix of Hamming codes, a lower bound of $OC(p_{Y|A,H_v}, p_{Y|A,H_w})$ asymptotically satisfies

$$OC(p_{Y|A,H_v}, p_{Y|A,H_w}) = \frac{1}{6}\log\left(\frac{\sigma^2+s(A)+1}{2\sqrt{\sigma^2+s(A)}}\right) \sim \frac{1}{6}\log\frac{\sigma}{2}.$$

So in the end, the minimum Chernoff information between any pair of hypotheses, under mixed measurements using Hamming codes as sensing vectors, is bigger than the Chernoff information obtained using separate observations. This means that mixed observations can offer strict improvement in the error exponent of hypothesis testing problems.

V. Characterization of Optimal Sensing Vector Maximizing the Error Exponent

In this section, we derive a characterization of the optimal deterministic time-varying measurements which maximize the error exponent of hypothesis testing. We further explicitly design the optimal measurements for some simple examples. We begin with the following lemma about the error exponent of hypothesis testing.

Lemma V.1. *Suppose that there are overall $l = \binom{n}{k}$ hypotheses. For any fixed k and n, the error exponent of the error probability of hypothesis testing is given by*

$$E = \min_{1\le v,w\le l, v\ne w} OC(p_{Y|A,H_v}, p_{Y|A,H_w}).$$

Proof: We first give an upper bound on the error probability of hypothesis testing. Without loss of generality, we assume H_1 is the true hypothesis. The error probability P_e in the Neyman-Pearson testing is stated as follows:

$$P_e = 2^{-mOC(p_{Y|A,H_v}, p_{Y|A,H_w})} \le 2^{-mE}.$$

By a union bound over the $l-1$ possible pairs (H_1, H_w), the probability that H_1 is not correctly identified as the true hypothesis is upper bounded by $l \times 2^{-mE}$ in terms of scaling.

Now we give a lower bound on the error probability of hypothesis testing. Without loss of generality, we assume that E is achieved between the hypothesis H_1 and the hypothesis H_2, namely,

$$E = OC(p_{Y|A,H_1}, p_{Y|A,H_2}).$$

Suppose that we are given the prior information that either hypothesis H_1 or H_2 is true. Knowing this prior information will not increase the error probability. Under this prior information, the error probability behaves asymptotically as 2^{-mE} as $m \to \infty$. This shows that the error exponent of hypothesis testing is no bigger than E. ■

The following theorem gives a simple characterization of the optimal probability density function $p_A(a)$ for the sensing vector. This enables us to explicitly find the optimal sensing vectors, under certain special cases of Gaussian random variables.

Theorem V.2. *In order to maximize the error exponent in hypothesis testing, the optimal sensing vectors have a distribution $p_A^*(a)$ which maximizes the minimum of the pairwise outer Chernoff information between different hypotheses:*

$$p_A^*(a) = \arg\max_{p_A(a)} \min_{1 \le v,w \le l, v \ne w} OC(p_{Y|A,H_v}, p_{Y|A,H_w}).$$

When $k = 1$ and the n random variables of interest are independent Gaussian random variables with the same variances, the optimal $p_A^*(a)$ admits a discrete probability distribution:

$$p_A^*(a) = \sum_{\sigma \text{ as a permutation}} \frac{1}{n!} \delta(a - \sigma(a^*)),$$

where a^* is a constant n-dimensional vector such that

$$a^* = \arg\max_{a} \sum_{1 \le v,w \le l, v \ne w} C(p_{Y|A,H_v}, p_{Y|A,H_w}).$$

Proof: The first statement follows from Lemma V.1. So we only need to prove the optimal sensing vectors for Gaussian random variables with the same variance under $k = 1$. We let $l = \binom{n}{k}$. For any $p_A(a)$, we have

$$\begin{aligned}
&\min_{1\le v,w\le l, v\neq w} OC(p_{Y|A,H_v}, p_{Y|A,H_w}) \\
\le\ & \frac{1}{\binom{n}{2}} \sum_{1\le v,w\le l, v\neq w} OC(p_{Y|A,H_v}, p_{Y|A,H_w}) \\
\le\ & \frac{1}{\binom{n}{2}} \sum_{1\le v,w\le l, v\neq w} \int p_A(a) C(p_{Y|A,H_v}, p_{Y|A,H_w})\, da \\
=\ & \frac{1}{\binom{n}{2}} \int p_A(a) \sum_{1\le v,w\le l, v\neq w} C(p_{Y|A,H_v}, p_{Y|A,H_w})\, da \\
\le\ & \frac{1}{\binom{n}{2}} \sum_{1\le v,w\le l, v\neq w} C(p_{Y|A^*,H_v}, p_{Y|A^*,H_w}),
\end{aligned}$$

where the second inequality follows from the fact that for any functions $f_i(\lambda)$, $i = 1, ..., l$, and for any λ in the intersection domain of all the $f_i(\lambda)$'s, $\min_\lambda(f_1(\lambda) + f_2(\lambda) + ... + f_l(\lambda)) \ge \min_\lambda f_1(\lambda) + \min_\lambda f_2(\lambda) + ... + \min_\lambda f_l(\lambda)$, and A^* is the same expression for a^* in the form of random variable.

On the other hand, for two Gaussian distributions with the same variances, the optimal λ in (VIII.1) is always equal to $\frac{1}{2}$, no matter what $p_A(a)$ is chosen. By symmetry, when

$$p_A^*(a) = \sum_{\sigma \text{ as a permutation}} \frac{1}{n!} \delta(a - \sigma(a^*)),$$

and $\lambda = \frac{1}{2}$, for any two different hypotheses H_v and H_w,

$$\begin{aligned}
& OC(p_{Y|A,H_v}, p_{Y|A,H_w}) \\
=\ & \int p_A^*(a) C(p_{Y|A,H_v}, p_{Y|A,H_w})\, da \\
=\ & \frac{1}{\binom{n}{2}} \int p_A^*(a) \sum_{1\le v,w\le l, v\neq w} C(p_{Y|A,H_v}, p_{Y|A,H_w})\, da \\
=\ & \int p_A^*(a) \frac{1}{\binom{n}{2}} \sum_{1\le v,w\le l, v\neq w} C(p_{Y|A^*,H_v}, p_{Y|A^*,H_w})\, da \\
=\ & \frac{1}{\binom{n}{2}} \sum_{1\le v,w\le l, v\neq w} C(p_{Y|A^*,H_v}, p_{Y|A^*,H_w}),
\end{aligned}$$

where the first equality is from the fact that $\lambda = \frac{1}{2}$ is the common maximizer for Chernoff information, the second equality is from the permutation symmetry of $p_A^*(a)$, the third equality is again from the generation of $p_A^*(a)$ from permutations of a^*, and the last equality follows from $\int p_A^*(a)\, da = 1$.

This means that under $p_A^*(a)$,

$$\begin{aligned}&\min_{1\le v,w\le l,v\neq w} OC(p_{Y|A,H_v}, p_{Y|A,H_w})\\ =\ &\frac{1}{\binom{n}{2}}\sum_{1\le v,w\le l,v\neq w} C(p_{Y|A^*,H_v}, p_{Y|A^*,H_w}).\end{aligned}$$

This further implies that the upper bound on $\min_{1\le v,w\le l,v\neq w} OC(p_{Y|A,H_v}, p_{Y|A,H_w})$ is achieved, and we can conclude that $p_A^*(a)$ is the optimal distribution. ■

We can use Theorem V.2 to calculate explicitly the optimal sensing vectors for n independent Gaussian random variables of the same variance σ^2, among which $k = 1$ random variable has a mean μ_1 and the mean of the other $(n-1)$ random variables is μ_2. To obtain the optimal measurements maximizing the error exponent E, we first need to calculate a constant vector a^* such that

$$a^* = \arg\max_a \sum_{1\le v,w\le l,v\neq w} C(p_{Y|A,H_v}, p_{Y|A,H_w}).$$

After a simple calculation, a^* is the optimal solution to

$$\begin{aligned}\underset{a}{\text{maximize}}\quad &\frac{1}{\binom{n}{2}}\sum_{1\le i,j\le n,i\neq j}(a_i-a_j)^2\\ \text{subject to}\quad &\sum_{i=1}^n a_i^2 \le 1,\end{aligned} \tag{V.1}$$

where a_i represents the i-th element of vector a, and the corresponding optimal error exponent is

$$\frac{1}{\binom{n}{2}}\sum_{1\le i,j\le n,i\neq j}\frac{(a_i^*-a_j^*)^2(\mu_1-\mu_2)^2}{8\sigma^2\sum_{i=1}^n (a_i^*)^2}.$$

This optimization problem is not a convex optimization program, however, it still admits zero duality gap from the S-procedure, and can be efficiently solved [25]. In fact, we get the optimal solution $a^* = [\frac{1}{\sqrt{2}}, \frac{-1}{\sqrt{2}}, 0, ..., 0]^T$. Then an optimal distribution $p_A^*(a)$ for the sensing vector is

$$p_A^*(a) = \sum_{\sigma \text{ as a permutation}} \frac{1}{n!}\delta(a-\sigma(a^*)).$$

Namely, this optimal sensing vector is to uniformly choose two random variables, say X_1 and X_2, and take their weighted sum $\frac{1}{\sqrt{2}}X_1 - \frac{1}{\sqrt{2}}X_2$. Correspondingly, the optimal error exponent is

$$\frac{(\mu_1-\mu_2)^2}{8\sigma^2}\frac{n}{\binom{n}{2}} = \frac{(\mu_1-\mu_2)^2}{4\sigma^2(n-1)}.$$

In contrast, if we perform separate observations of n random variables individually, the error exponent will be

$$\frac{(\mu_1-\mu_2)^2}{8\sigma^2}\frac{n-1}{\binom{n}{2}}=\frac{(\mu_1-\mu_2)^2}{4\sigma^2 n}.$$

In fact, linear mixed observations increase the error exponent by $\frac{n}{n-1}$ times. When n is small, the improvement is significant. For example, when $n=2$, the error exponent is doubled.

From this example of Gaussian random variables with different means, we can make some interesting observations. On the one hand, quite surprisingly, separate observations of random variables are *not* optimal in maximizing the error exponent. On the other hand, in the optimal measurement scheme, each measurement takes a linear mixing of only two random variables, instead of mixing all the random variables. It is interesting to see whether these observations hold for other types of random variables.

VI. Efficient Algorithms for Hypothesis Testing from Mixed Observations

In Section III, we introduce the likelihood ratio test algorithms for three different types of mixed observations. Even though the likelihood ratio test is the optimal test achieving the smallest hypothesis testing error probability, conducting the likelihood ratio test over $\binom{n}{k}$ hypotheses is computationally challenging. Especially, when n and k are large, finding k abnormal random variables out of n random variables is almost impossible by using the likelihood ratio test. To overcome this drawback, we further design efficient algorithms to find k abnormal random variables among n random variables for large n and k by using Least Absolute Shrinkage and Selection Operator (LASSO) algorithm [26], and Message Passing (MP) algorithm [27].

A. LASSO based hypothesis testing algorithm

We propose to use the LASSO algorithm to detect k anomalous random variables when they have different means from the other $n-k$ random variables. Without loss of generality, we assume that each of the abnormal random variables has a non-zero mean, while each of the other $(n-k)$ random variables has mean 0.

LASSO, also known as Basis Pursuit denoising [28], is a well-known sparse regression tool in statistics, and has been successfully applied to various fields such as signal processing [29],

machine learning [30], and control system [31]. For a sensing matrix $A \in R^{m \times n}$, and observation $y \in R^m$, the LASSO optimization formulation is given as follows:

$$\underset{x}{\text{minimize}} \ \frac{1}{2}\|y - Ax\|_2^2 + \lambda\|x\|_1, \tag{VI.1}$$

where x is a variable, and λ is a parameter for the penalty term, ℓ_1 norm of x, which makes more elements of x driven to zero as λ increases.

We use the LASSO formulation (VI.1) to obtain a solution $\hat{x}$, by taking $y = [Y^1, ..., Y^m]^T$ and A as an $m \times n$ matrix with its j-th row equal to the sensing vector A^j. Since we are interested in finding k abnormal random variables, we solve (VI.1) and select k largest elements of $\hat{x}$ in amplitude. We decide the corresponding k random variables as the k abnormal random variables.

B. MP based hypothesis testing algorithm

We further design a message passing method to discover general abnormal random variables, even if the abnormal random variables have the same mean as the regular random variables. Our message passing algorithm uses *bipartite graph* to perform statistical inference. We remark that message passing algorithms have been successfully applied to decoding for error correcting code [27], including Low Density Parity Check (LDPC) code.

Let us denote the j-th linear mixed function by ${a^j}^T X$, where X is a random variable vector, and a^j is the j-th sensing vector for a linear mixed observation. The observation random variable at time j, i.e. Y^j, is simply represented as ${a^j}^T X^j$ as in (II.2). Note that over time indices j's, the random variable vector X takes independent realizations. Now we define a bipartite factor graph as follows. We will represent a random variable X_i using a *variable node* on the left, and represent an observation Y^j as a *check node* on the right. Variable nodes and check nodes appear on two sides of the bipartite factor graph. A line is drawn between a node X_i and a node Y^j if and only if that random variable X_i is involved in the observation Y^j. We call the observation Y^j linked with X_i as a *neighbor* of X_i. We will denote the set of neighbors of X_i by $N(X_i)$. Similarly, the set of random variable X_i's linked with a random variable Y^j is denoted by $N(Y^j)$. Fig. 1(a) is an example of the factor graph for a given sensing matrix A in Fig. 1(b). In our message passing algorithm, messages are exchanged between variable nodes and check nodes.

The messages exchanged between variable nodes and check nodes are the probabilities that a random variable X_i is abnormal. More precisely, the message sent from a variable node

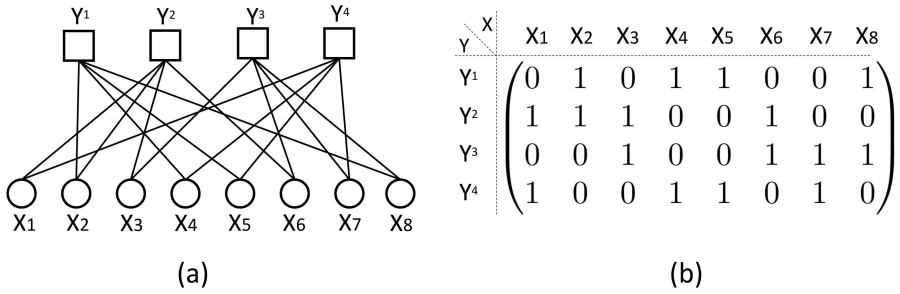

Fig. 1. Illustration of a factor graph (a) from a matrix (b). A random variable X_i and Y^j are considered as a variable node and a check node in the graph respectively.

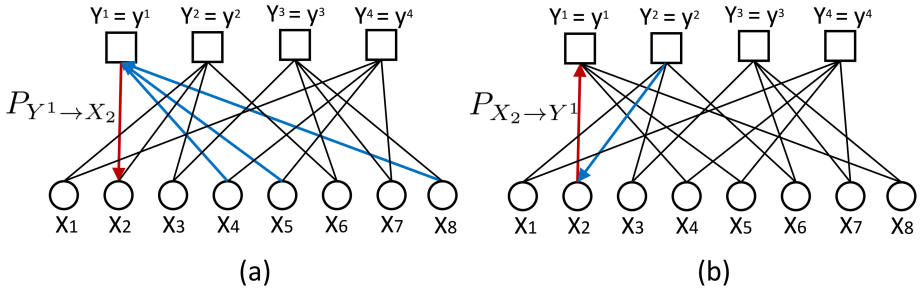

Fig. 2. (a) Message sent from a check node to a variable node. The message sent from a check node Y^1 to a variable node X_2 (red arrow) is expressed as the probability $P_{Y^1 \to X_2}(X_2 \text{ is abnormal} \mid Y^1 = y^1)$ by considering probabilities $P_{X_i \to Y^1}$, $i = 4, 5, 8$ (blue arrow). (b) Message sent from a variable node to a check node. The message sent from a variable node X_2 to a check node Y^1 (red arrow) is expressed as the probability $P_{X_2 \to Y^1}(X_2 \text{ is abnormal})$ by considering probability $P_{Y^2 \to X_2}$.

X_i to a check node Y^j is the probability that the variable node X_i is abnormal, namely $P_{X_i \to Y^j}(X_i \text{ is abnormal})$, based on local information gathered at variable node X_i. Similarly, the message sent from a check node Y^j to a variable node X_i is the probability that the variable node X_i is abnormal based on local information gathered at variable node Y^j. For example, Fig. 2(a) shows the message sent from a check node Y^1 to a variable node X_2, which is the probability that X_2 is abnormal given observation y^1 and incoming messages from X_4 to Y^1, X_5 to Y^1, and X_8 to Y^1. Fig. 2(b) illustrates the message from a variable node X_2 to a check node Y_1, which is the probability that X_2 is abnormal when we consider the incoming message from Y^2 to X_2. The message $\mathfrak{m}$ from a check node Y^j to a variable node X_i is expressed as a

function of incoming messages from neighbors of X_i and Y^j, and given by

$$\begin{aligned}\mathfrak{m}_{Y^j\to X_i} &= f\big(\mathfrak{m}_{X_{l\in N(Y^j)\setminus X_i}\to Y^j}, Y^j\big)\\ &= P_{Y_j\to X_i}\big(X_i \text{ is abnormal} \mid Y_j = y_j\big),\end{aligned} \tag{VI.2}$$

where $f(\cdot)$ is a function calculating the probability of X_i being abnormal based on incoming messages from its neighbor variable nodes and realized observation y^j. The message m from a variable node to a check node is expressed by

$$\begin{aligned}\mathfrak{m}_{X_i\to Y^j} &= h\big(\mathfrak{m}_{Y^{l\in N(X_j)\setminus Y^j}\to X_i}\big)\\ &= P_{X_i\to Y^j}\big(X_i \text{ is abnormal}\big),\end{aligned} \tag{VI.3}$$

where $h(\cdot)$ is a function calculating the probability of X_i being abnormal based on incoming messages from check nodes. In the same way, we calculate the probability that X_i is normal.

VII. Numerical Experiments

In this section, we numerically evaluate the performance of mixed observations in hypothesis testing. We first simulate the error probability of identifying anomalous random variables through linear mixed observations. The linear mixing used in the simulation is based on sparse bipartite graphs. In sparse bipartite graphs [32–34], n variable nodes on the left are used to represent the n random variables, and m measurement nodes on the right are used to represent the m measurements. If and only if the i-th random variable is nontrivially involved in the j-th measurement, there is an edge connecting the i-th variable node to the j-th measurement node. Unlike sparse bipartite graphs already used in LDPC codes, and compressed sensing [32–34], a novelty in this paper is that our sparse bipartite graphs are allowed to have more measurement nodes than variable nodes, namely $m \geq n$. In this simulation, there are 6 edges emanating from each measurement node on the right, and there are $\frac{6m}{n}$ edges emanating from each variable node on the left. After a uniformly random permutation, the $6m$ edges emanating from the measurement nodes are plugged into the $6m$ edge "sockets" of the left variable nodes. If there is an edge connecting the i-th variable node to the j-th measurement node, then the linear mixing coefficient before the i-th random variable in the j-th measurement is set to 1; otherwise that linear mixing coefficient is set to 0. We call the Likelihood Ratio Test (LRT) with separate observations and mixed (compressed) observations as SLRT and CLRT respectively. We also abbreviate the message passing and LASSO based hypothesis testing method to MP and LASSO in the figures below.

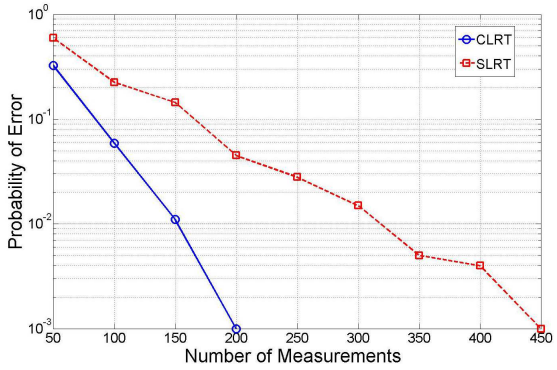

Fig. 3. Error probability as a function of m in log scale when $(n, k) = (100, 1)$. The normal and abnormal random variables follow $\mathcal{N}(0, 1)$ and $\mathcal{N}(0, 100)$ respectively.

A. Random variables with different variances

In the first simulation example, we take $n = 100$, and let m vary from 50 to 450. We assume that $k = 1$ random variable follows the Gaussian distribution $\mathcal{N}(0, 100)$, and the other $(n-k) = 99$ random variables follow another distribution $\mathcal{N}(0, 1)$. We used likelihood ratio test algorithms to find the anomalous random variables through the described linear mixed observations based on sparse bipartite graphs. For comparison, we also implement the likelihood ratio test algorithms for separate observations of random variables, where we first make $\lfloor \frac{m}{n} \rfloor$ separate observations of each random variables, and then made an additional separate observation for uniformly randomly selected $(m \mod n)$ random variables. For each m, we perform 1000 random trials, and record the number of trials failing to identify the anomalous random variables. The error probability, as a function of m, is plotted in Fig. 3. We can see that linear mixed observation offers significant reduction in the error probability of hypothesis testing, under the same number of observations.

We further carry out simulations for $n = 200$. Fig. 4 and 5 show the hypothesis testing error probability when normal and abnormal random variables follow the Gaussian distributions $\mathcal{N}(0, 1)$ and $\mathcal{N}(0, 100)$ on $k = 1$ and 2 respectively. The hypothesis testing error probability is obtained from 1000 random trials. In this simulation, we use MP based hypothesis testing algorithm and compare it against LRT methods.

For large values of n and k, Fig. 6 shows the performance of MP based hypothesis testing algorithm when the two types of random variables have different variances. We set k from 1 to 11 when $n = 1000$. In this parameter setup, LRT methods have difficulties in finding the k

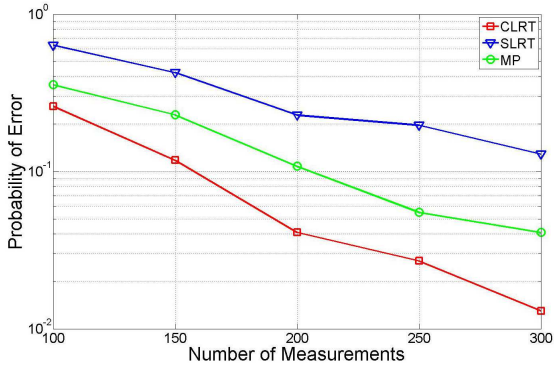

Fig. 4. Error probability as a function of m in log scale when $(n, k) = (200, 1)$. The normal and abnormal random variables follow $\mathcal{N}(0, 1)$ and $\mathcal{N}(0, 100)$ respectively.

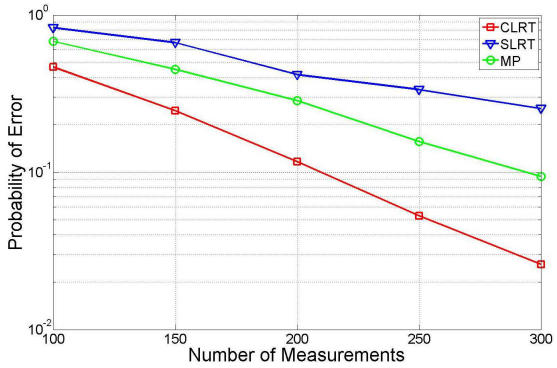

Fig. 5. Error probability as a function of m in log scale when $(n, k) = (200, 2)$. The normal and abnormal random variables follow $\mathcal{N}(0, 1)$ and $\mathcal{N}(0, 100)$ respectively.

abnormal random variables out of n. The normal and abnormal random variables follow the Gaussian distribution $\mathcal{N}(0, 1)$ and $\mathcal{N}(0, 100)$ respectively, and the error probability is obtained from 500 random trials.

B. Random variables with different means

Under the same simulation setup as in Fig. 3, we test the error probability performance of mixed observations for two Gaussian distributions: the anomalous Gaussian distribution $\mathcal{N}(0, 1)$, and the normal Gaussian distribution $\mathcal{N}(8, 1)$. We also slightly adjust the number of total random variables as $n = 102$, to make sure that each random variable participates in the same integer number $\frac{6m}{n}$ of measurements. Mixed observations visibly reduce the error probability under the same number of measurements, compared with separate observations. For example, even when

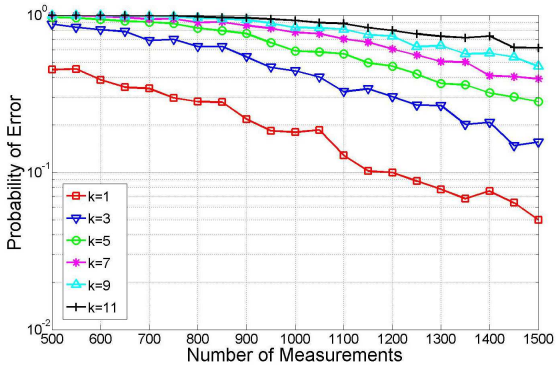

Fig. 6. Error probability of MP based hypothesis testing algorithm as a function of m in log scale on various k values when n = 1000. The normal and abnormal random variables follow $\mathcal{N}(0,1)$ and $\mathcal{N}(0,100)$ respectively.

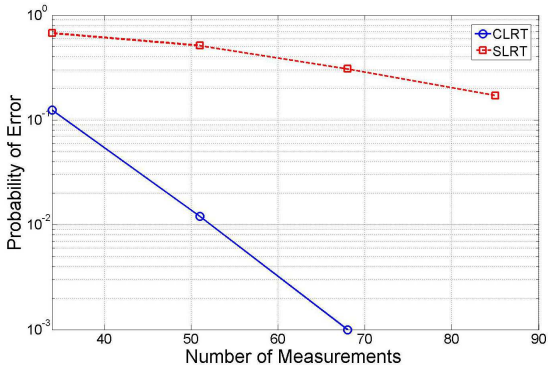

Fig. 7. Error probability as a function of m in log scale when n = 102. The normal and abnormal random variables follow $\mathcal{N}(8,1)$ and $\mathcal{N}(0,1)$ respectively.

$m = 68 < n = 102$, CLRT correctly identifies the anomalous random variable in 999 out of 1000 cases by using m mixed observations from the bipartite graphs. Fig. 7 shows the result when the two types of random variables have different expectations.

In addition, we carry out simulations to show the results from MP and LASSO based hypothesis testing methods and compare them against the results from LRT methods. Fig. 8 and 9 have the simulation results when the normal and abnormal random variables follow the Gaussian distribution $\mathcal{N}(0,1)$ and $\mathcal{N}(7,1)$ on $k = 1$ and 2 respectively. The error probability is obtained from 1000 random trials.

We compare the performance of MP and LASSO based hypothesis testing algorithms on large values of n and k which is the computational challenging case for LRT in Fig. 10 and 11. For this simulation, we set n to 1000 and k from 1 to 11, and for each m, we perform 500 random

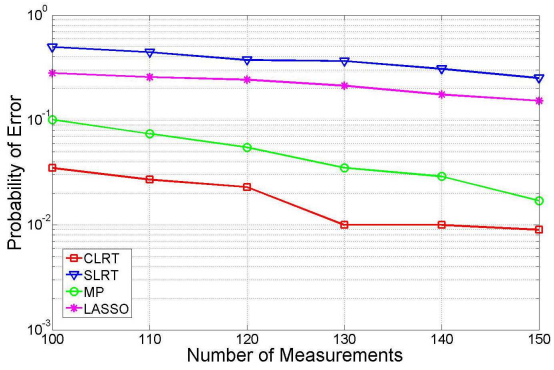

Fig. 8. Error probability as a function of m in log scale when $(n, k) = (200, 1)$. The normal and abnormal random variables follow $\mathcal{N}(0, 1)$ and $\mathcal{N}(7, 1)$ respectively.

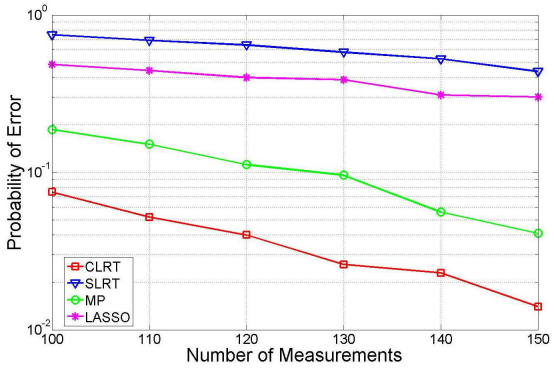

Fig. 9. Error probability as a function of m in log scale when $(n, k) = (200, 2)$. The normal and abnormal random variables follow $\mathcal{N}(0, 1)$ and $\mathcal{N}(7, 1)$ respectively.

trials to obtain the error probability. The normal and abnormal random variables follow $\mathcal{N}(0, 1)$ and $\mathcal{N}(8, 1)$ respectively. Finally, Fig. 12 shows the comparison result between MP and LASSO based hypothesis testing methods.

C. Random variables with different means and variances

We further test the error probability performance of mixed observations for two Gaussian distributions with different means and variances. Fig. 13 and 14 show the results when two types of random variables have different means and variances. The abnormal and normal random variables follow $\mathcal{N}(0, 1)$, and $\mathcal{N}(7, 100)$ respectively. We obtain the error probability from 1000 random trials. The simulation results show that CLRT and MP still identify the abnormal random variables with fewer observations than SLRT.

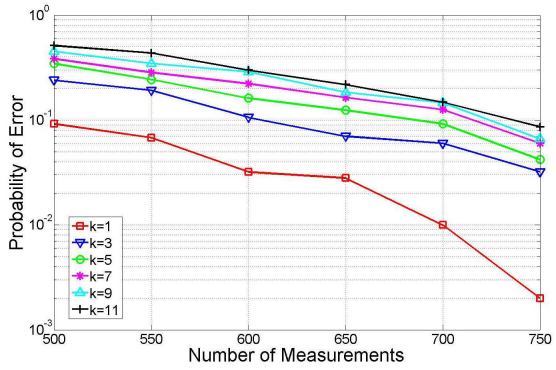

Fig. 10. Error probability of MP based hypothesis testing algorithm as a function of m in log scale on various k values when n = 1000. The normal and abnormal random variables follow $\mathcal{N}(0,1)$ and $\mathcal{N}(8,1)$ respectively.

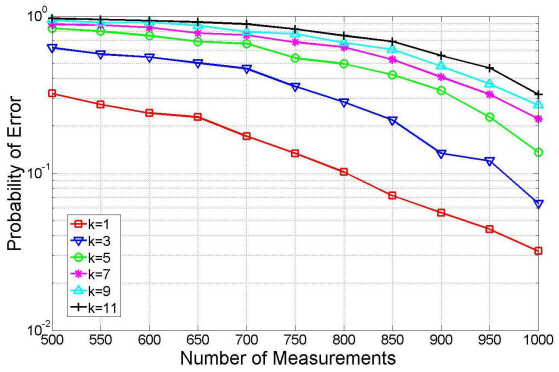

Fig. 11. Error probability of LASSO based hypothesist testing algorithm as a function of m in log scale on various k values when n = 1000. The normal and abnormal random variables follow $\mathcal{N}(0,1)$ and $\mathcal{N}(8,1)$ respectively.

VIII. Conclusion

In this paper, we have studied the problem finding k anomalous random variables following a different probability distribution among n random variables, by using *non-adaptive* mixed observations of these n random variables. Our analysis has shown that mixed observations, compared with separate observations of individual random variables, can reduce the number of samples required to identify the anomalous random variables accurately. Compared with conventional compressed sensing problems, in our setting, each random variable may take dramatically different realizations in different observations.

There are some questions that remain open in performing hypothesis testing from mixed observations. For example, for random variables of non-Gaussian distributions, it is not explicitly known what linear mixed observations maximize the anomaly detection error exponents. In

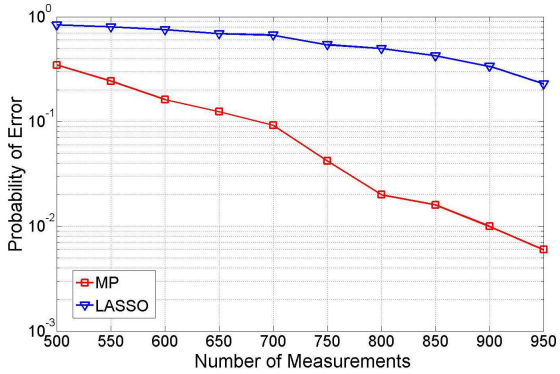

Fig. 12. Error probability as a function of m in log scale on various k values when $(n,k) = (1000,5)$. The normal and abnormal random variables follow $\mathcal{N}(0,1)$ and $\mathcal{N}(8,1)$ respectively.

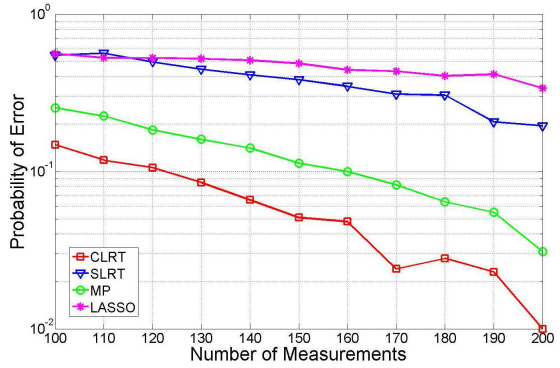

Fig. 13. Error probability as a function of m in log scale when $(n,k) = (200,1)$. The normal and abnormal random variables follow $\mathcal{N}(0,1)$ and $\mathcal{N}(7,100)$ respectively.

addition, it is very interesting to explore the mixed observations for anomaly detection for random variables with unknown abnormal probability distributions.

APPENDIX: PROOF OF THEOREM III.3

We provide the proof of Theorem III.3 here. The framework of this proof follows the book of Cover [23, Chap. 11]

Proof: In Algorithm 6, for two different hypotheses H_v and H_w, we choose the probability likelihood ratio threshold of the Neyman-Pearson testing in a way, such that the hypothesis testing error probability decreases with the largest error exponent. Now we focus on deriving what this largest error exponent is, under deterministic time-varying measurements. .

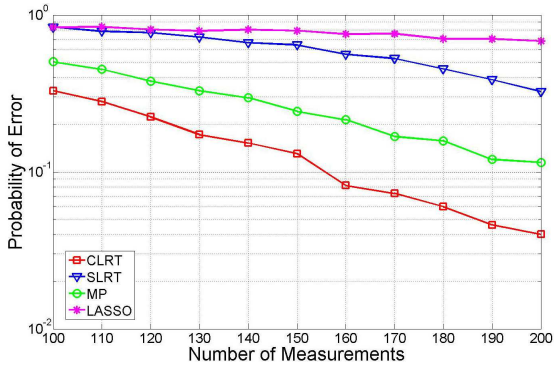

Fig. 14. Error probability as a function of m in log scale when $(n, k) = (200, 2)$. The normal and abnormal random variables follow $\mathcal{N}(0, 1)$ and $\mathcal{N}(7, 100)$ respectively.

For simplicity of presentation, we first consider a special case: there are only two possible sensing vectors a_1 and a_2; and one half of the sensing vectors are a_1 while the other half are a_2. The conclusions can be extended to general distribution $p_A(a)$ on A, in a similar way of reasoning. In addition, we assume that the observation data is over a discrete space χ, which can also be generalized to a continuous space without affecting the conclusion in this theorem. Since we use the probability mass function over a discrete space, we will use upper letters for the probability mass function to distinguish the probability density function over a continuous space in this proof. Suppose we take m measurements in total, our assumption translates to that $\frac{1}{2}m$ measurements are taken from the sensing vector a_1, and $\frac{1}{2}m$ measurements are taken from the sensing vector a_2. Without loss of generality, we consider two hypotheses denoted by H_1 and H_2. Under the sensing vector a_1, we assume that H_1 generates distribution P_1 for observation data; H_2 generates distribution P_2 for observation data. Under a_2, we assume that H_1 generates distribution P_3 for observation data; H_2 generates distribution P_4 for observation data. Please refer to Table II for the observation distributions under different sensing vectors and different hypotheses.

Suppose that P is the empirical distribution of observation data under the sensing vector a_1, and that P' is the empirical distribution of observation data under the sensing vector a_2. Then the Neyman-Pearson testing decides that hypothesis H_1 is true if, for a certain constant T,

$$\frac{1}{2}[D(P||P_2) - D(P||P_1)] + \frac{1}{2}[D(P'||P_4) - D(P'||P_3)] \geq \frac{1}{n}\log(T),$$

where $D(P||Q) \triangleq \sum_{y\in\mathcal{X}} P(y)\log\frac{P(y)}{Q(y)}$, which is the relative entropy or Kullback-Leibler distance between two probability mass functions P and Q. By the Sanov's theorem [23], the error

TABLE II

HYPOTHESIS TESTING MEASUREMENT DISTRIBUTION

	H_1	H_2
a_1	P_1	P_2
a_2	P_3	P_4

exponent of the second kind, namely wrongly deciding "hypothesis H_1 is true" when hypothesis H_2 is actually true, is given by the following optimization problem:

$$\begin{aligned}
&\underset{P,P'}{\text{minimize}} \ \ \frac{1}{2}D(P\|P_2)+\frac{1}{2}D(P'\|P_4)\\
&\text{subject to}\\
&\frac{1}{2}[D(P\|P_2)-D(P\|P_1)]+\frac{1}{2}[D(P'\|P_4)-D(P'\|P_3)]\geq\frac{\log(T)}{n},\\
&\sum_y P(y)=1,\\
&\sum_y P'(y)=1.
\end{aligned}$$

Using the Lagrange multiplier method, we try to minimize

$$\begin{aligned}
&L(P,P',\lambda,v_1,v_2)\\
&=\frac{D(P\|P_2)}{2}+\frac{D(P'\|P_4)}{2}\\
&\quad+\lambda\left(\frac{D(P\|P_2)-D(P\|P_1)}{2}+\frac{D(P'\|P_4)-D(P'\|P_3)}{2}-\frac{\log(T)}{n}\right)\\
&\quad+v_1\left(\sum_y P(y)-1\right)+v_2\left(\sum_y P'(y)-1\right).
\end{aligned}$$

Differentiating with respect to $P(y)$ and $P'(y)$, we get

$$\begin{aligned}
\frac{1}{2}[\log(\frac{P(y)}{P_2(y)})+1+\lambda\log(\frac{P_1(y)}{P_2(y)})]+v_1=0,\\
\frac{1}{2}[\log(\frac{P'(y)}{P_4(y)})+1+\lambda\log(\frac{P_3(y)}{P_4(y)})]+v_2=0.
\end{aligned}$$

From these equations, we can obtain the minimizing P and P',

$$\begin{aligned}
P=P_\lambda(y|a_1)=\frac{P_1^\lambda(y)P_2^{1-\lambda}(y)}{\sum_{y\epsilon\chi}P_1^\lambda(y)P_2^{1-\lambda}(y)},\\
P'=P_\lambda(y|a_2)=\frac{P_3^\lambda(y)P_4^{1-\lambda}(y)}{\sum_{y\epsilon\chi}P_3^\lambda(y)P_4^{1-\lambda}(y)},
\end{aligned}$$

where λ is chosen such that $\frac{1}{2}[D(P\|P_2) - D(P\|P_1)] + \frac{1}{2}[D(P'\|P_4) - D(P'\|P_3)] = \frac{1}{n}\log(T)$.

By symmetry, the error exponent of the second kind, and the error exponent of the first kind are stated as follows respectively:

$$\frac{1}{2}D(P_\lambda(y|a_1) \parallel P_2) + \frac{1}{2}D(P_\lambda(y|a_2) \parallel P_4),$$
$$\frac{1}{2}D(P_\lambda(y|a_1) \parallel P_1) + \frac{1}{2}D(P_\lambda(y|a_2) \parallel P_3).$$

The first exponent is a non-decreasing function in λ, and the second exponent is a non-increasing function in λ. In fact, the optimal error exponent, which is the minimum of these two exponents, is achieved when they are equal:

$$\begin{aligned}&\frac{1}{2}D(P_\lambda(y|a_1) \parallel P_2) + \frac{1}{2}D(P_\lambda(y|a_2) \parallel P_4)\\ =\ &\frac{1}{2}D(P_\lambda(y|a_1) \parallel P_1) + \frac{1}{2}D(P_\lambda(y|a_2) \parallel P_3).\end{aligned}$$

This finishes the characterization of the optimal error exponent in pairwise hypothesis testing, under deterministic time-varying measurements.

We further prove that this error exponent is equivalent to

$$\begin{aligned}OC(P_{Y|A,H_v}, P_{Y|A,H_w}) &= -\min_{0\le\lambda\le1} \int P_A(a)\log\left(\int P^{\lambda}_{Y|A,H_v}(y)P^{1-\lambda}_{Y|A,H_w}(y)\,dy\right)da\\ &= -\min_{0\le\lambda\le1} E_A\left(\log\left(\int P^{\lambda}_{Y|A,H_v}(y)P^{1-\lambda}_{Y|A,H_w}(y)\,dy\right)\right). \qquad \text{(VIII.1)}\end{aligned}$$

In the proof, we restrict our attention to $H_v = H_1$ and $H_w = H_2$. We will show that the $\lambda \in [0,1]$ that minimizes $E_A\left(\log\left(\int P^{\lambda}_{Y|A,H_v}(y)P^{1-\lambda}_{Y|A,H_w}(y)\,dy\right)\right)$ exactly leads to the following equality:

$$\begin{aligned}&\frac{1}{2}D(P_\lambda(y|a_1) \parallel P_2) + \frac{1}{2}D(P_\lambda(y|a_2) \parallel P_4)\\ =\ &\frac{1}{2}D(P_\lambda(y|a_1) \parallel P_1) + \frac{1}{2}D(P_\lambda(y|a_2) \parallel P_3). \qquad \text{(VIII.2)}\end{aligned}$$

Moreover, under that minimizer λ, we also obtain the following equality:

$$\begin{aligned}&-E_A\left(\log\left(\int P^{\lambda}_{Y|A,H_v}(y)P^{1-\lambda}_{Y|A,H_w}(y)\,dy\right)\right)\\ =\ &\frac{1}{2}D(P_\lambda(y|a_1) \parallel P_1) + \frac{1}{2}D(P_\lambda(y|a_2) \parallel P_3). \qquad \text{(VIII.3)}\end{aligned}$$

On the one hand, we obtain following equations from (VIII.2) and the definition of the relative entropy:

$$\begin{aligned}0 &\overset{\text{(VIII.2)}}{=} \frac{1}{2}[D(P_\lambda(y|a_1) \parallel P_1) - D(P_\lambda(y|a_1) \parallel P_2)] + \frac{1}{2}[D(P_\lambda(y|a_2) \parallel P_3) - D(P_\lambda(y|a_2) \parallel P_4)]\\ &= \frac{1}{2}\frac{\sum_y P_1^{\lambda}(y)P_2^{1-\lambda}(y)\log(\frac{P_2(y)}{P_1(y)})}{\sum_y P_1^{\lambda}(y)P_2^{1-\lambda}(y)} + \frac{1}{2}\frac{\sum_y P_3^{\lambda}(y)P_4^{1-\lambda}(y)\log(\frac{P_4(y)}{P_3(y)})}{\sum_y P_3^{\lambda}(y)P_4^{1-\lambda}(y)}. \qquad \text{(VIII.4)}\end{aligned}$$

Let a particular λ satisfy (VIII.4). Under this λ, the hypothesis testing error exponent is equal to

$$\frac{1}{2}D(P_\lambda(y|a_1) \parallel P_1) + \frac{1}{2}D(P_\lambda(y|a_2) \parallel P_3),$$

and it is further derived from the definition of the relative entropy as follows:

$$\begin{aligned}
&\frac{1}{2}\frac{1}{\sum_y P_1^\lambda(y)P_2^{1-\lambda}(y)}\sum_y\Bigg(P_1^\lambda(y)P_2^{1-\lambda}(y) \\
&\qquad\qquad\left[(1-\lambda)\log(\frac{P_2(y)}{P_1(y)}) - \log\big(\sum_y P_1^\lambda(y)P_2^{1-\lambda}(y)\big)\right]\Bigg) \\
&+\frac{1}{2}\frac{1}{\sum_y P_3^\lambda(y)P_4^{1-\lambda}(y)}\sum_y\Bigg(P_3^\lambda(y)P_4^{1-\lambda}(y) \\
&\qquad\qquad\left[(1-\lambda)\log(\frac{P_4(y)}{P_3(y)}) - \log\big(\sum_y P_3^\lambda(y)P_4^{1-\lambda}(y)\big)\right]\Bigg).
\end{aligned}$$

Since the first parts of both summations are equal to 0 from (VIII.4), we have the following equation:

$$\begin{aligned}
&\frac{1}{2}D(P_\lambda(y|a_1) \parallel P_1) + \frac{1}{2}D(P_\lambda(y|a_2) \parallel P_3) \\
&= -\frac{1}{2}\log\big(\sum_y P_1^\lambda(y)P_2^{1-\lambda}(y)\big) - \frac{1}{2}\log\big(\sum_y P_3^\lambda(y)P_4^{1-\lambda}(y)\big),
\end{aligned} \tag{VIII.5}$$

which is just $-E_A\left(\log\left(\int P_{Y|A,H_v}^\lambda(y)P_{Y|A,H_w}^{1-\lambda}(y)\,dy\right)\right)$ under this λ achieving (VIII.2).

On the other hand, to minimize

$$E_A\left(\log\left(\int P_{Y|A,H_v}^\lambda(y)P_{Y|A,H_w}^{1-\lambda}(y)\,dy\right)\right),$$

we set the derivative of (VIII.5) with respect to λ to 0 as follows:

$$0 = \frac{1}{2}\frac{\sum_y P_1^\lambda(y)P_2^{1-\lambda}(y)\log(\frac{P_2(y)}{P_1(y)})}{\sum_y P_1^\lambda(y)P_2^{1-\lambda}(y)} + \frac{1}{2}\frac{\sum_y P_3^\lambda(y)P_4^{1-\lambda}(y)\log(\frac{P_4(y)}{P_3(y)})}{\sum_y P_3^\lambda(y)P_4^{1-\lambda}(y)}. \tag{VIII.6}$$

It is noteworthy that (VIII.6) is the same as (VIII.4). Let us denote a minimizer

$$\lambda_{min} = \arg\min_{0\le\lambda\le1} E_A\left(\log\left(\int P_{Y|A,H_v}^\lambda(y)P_{Y|A,H_w}^{1-\lambda}(y)\,dy\right)\right).$$

Then, when $\lambda = \lambda_{min}$, (VIII.4) is satisfied. Furthermore, for $\lambda = \lambda_{min}$, we obtain (VIII.5). It states that for $\lambda = \lambda_{min}$, the following equation is satisfied:

$$\begin{aligned}
&- E_A\left(\log\left(\int P_{Y|A,H_v}^\lambda(y)P_{Y|A,H_v}^{1-\lambda}(y)\,dy\right)\right) \\
&= \frac{1}{2}D(P_\lambda(y|a_1) \parallel P_1) + \frac{1}{2}D(P_\lambda(y|a_2) \parallel P_3).
\end{aligned} \tag{VIII.7}$$

Therefore, we conclude the equivalence of the two different definitions of outer Chernoff information.

Overall, the smallest possible error exponent between any pair of hypotheses is

$$E = \min_{1\le v,w\le l, v\neq w} OC(P_{Y|A,H_v}, P_{Y|A,H_w}).$$

Without loss of generality, we assume H_1 is the true hypothesis. Since the error probability P_e in the Neyman-Pearson testing is

$$P_e \triangleq 2^{-mOC(P_{Y|A,H_v},P_{Y|A,H_w})} \le 2^{-mE}.$$

By a union bound over the $l-1$ possible pairs (H_1, H_w), the probability that H_1 is not correctly identified as the true hypothesis is upper bounded by $l\times 2^{-mE}$ in terms of scaling, where $l = \binom{n}{k}$. Therefore, $m = \Theta(k\log(n)E^{-1})$ samplings are enough for identifying the k anomalous samples with high probability. When E grows polynomially with n, this implies a significant reduction in the number of samples needed. ■

REFERENCES

[1] H. V. Poor and O. Hadjiliadis, *Quickest Detection*. Cambridge University Press, 2008.

[2] N. Duffield and F. Lo Presti, "Multicast inference of packet delay variance at interior network links," in *Proceedings of the IEEE Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies*, 2000, pp. 1351–1360.

[3] D. Jaruskova, "Some problems with application of change-point detection methods to environmental data," *Environmetrics*, vol. 8, no. 5, pp. 469–483, 1997.

[4] V. F. Pisarenko, A. F. Kushnir, and I. V. Savin, "Statistical adaptive algorithms for estimation of onset moments of seismic phases," *Physics of the Earth and Planetary Interiors*, vol. 47, pp. 4–10, 1987.

[5] F. Lo Presti, N. G. Duffield, J. Horowitz, and D. Towsley, "Multicast-based inference of network-internal delay distributions," *IEEE/ACM Transactions on Networking*, vol. 10, no. 6, pp. 761–775, 2002.

[6] Y. Xia and D. Tse, "Inference of link delay in communication networks," *IEEE Journal on Selected Areas in Communications*, vol. 24, no. 12, pp. 2235–2248, 2006.

[7] L. Lai, Y. Fan, and H. V. Poor, "Quickest detection in cognitive radio: A sequential change detection framework," in *Proceedings of IEEE Global Telecommunications Conference (GLOBECOM)*, 2008, pp. 1–5.

[8] C. E. Cook and M. Bernfeld, *Radar signals: an introduction to theory and application*. Artech Print on Demand, 1993.

[9] M. Basseville and I. Nikiforov, *Detection of Abrupt Changes: Theory and Application*. Prentice-Hall, 1993.

[10] L. Lai, H. V. Poor, Y. Xin, and G. Georgiadis, "Quickest search over multiple sequences," *IEEE Transactions on Information Theory*, vol. 57, no. 8, pp. 5375–5386, 2011.

[11] M. Malloy and R. Nowak, "On the limits of sequential testing in high dimensions," in *Proceedings of Asilomar Conference on Signals, Systems and Computers*, 2011, pp. 1245–1249.

[12] ——, "Sequential analysis in high dimensional multiple testing and sparse recovery," in *Proceedings of IEEE International Symposium on Information Theory (ISIT)*, 2011, pp. 2661–2665.

[13] M. Malloy, G. Tang, and R. Nowak, "Quickest search for a rare distribution," in *Proceedings of the Annual Conference on Information Sciences and Systems (CISS)*, 2012, pp. 1–6.

[14] R. Caromi, Y. Xin, and L. Lai, "Fast multi-band spectrum scanning for cognitive radio systems," *IEEE Transactions on Communications*, vol. 61, no. 1, pp. 63–75, 2013.

[15] E. Candès and T. Tao, "Decoding by linear programming," *IEEE Transactions on Information Theory*, vol. 51, no. 12, pp. 4203–4215, 2005.

[16] D. L. Donoho, "High-dimensional centrally symmetric polytopes with neighborliness proportional to dimension," *Discrete & Computational Geometry*, vol. 35, no. 4, pp. 617–652, 2006.

[17] D. L. Donoho and J. Tanner, "Neighborliness of randomly projected simplices in high dimensions," in *Proceedings of the National Academy of Sciences of the United States of America*, 2005, pp. 9452–9457.

[18] S. Ji, Y. Xue, and L. Carin, "Bayesian compressive sensing," *IEEE Transactions on Signal Processing*, vol. 56, no. 6, pp. 2346–2356, 2008.

[19] G. Yu and G. Sapiro, "Statistical compressive sensing of gaussian mixture models," in *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2011, pp. 3728–3731.

[20] D. Angelosante, G. Giannakis, and E. Grossi, "Compressed sensing of time-varying signals," in *Proceedings of IEEE International Conference on Digital Signal Processing*, 2009, pp. 1–8.

[21] N. Vaswani, "Kalman filtered compressed sensing," in *Proceedings of IEEE International Conference on Image Processing*, 2008, pp. 893–896.

[22] E. Karseras and W. Dai, "Tracking dynamic sparse signals: a hierarchical kalman filter," in *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2013, pp. 6546–6550.

[23] T. Cover and J. Thomas, *Elements of Information Theory, 2nd Edition*. Wiley, 2006.

[24] F. Nielsen, "Chernoff information of exponential families," *http://arxiv.org/abs/1102.2684*, 2011.

[25] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.

[26] R. Tibshirani, "Regression shrinkage and selection via the lasso," *Journal of the Royal Statistical Society. Series B (Methodological)*, vol. 58, no. 1, pp. 267–288, 1996.

[27] D. J. C. MacKay, *Information theory, inference and learning algorithms*. Cambridge University Press, 2003.

[28] S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," *SIAM journal on scientific computing*, vol. 20, no. 1, pp. 33–61, 1998.

[29] S. Chen and D. L. Donoho, "Examples of basis pursuit," in *Proceedings of Wavelet Applications in Signal and Image Processing III*, vol. 2569, 1995, pp. 564–574.

[30] S. Sra, S. Nowozin, and S. J. Wright, *Optimization for Machine Learning*. The MIT Press, 2011.

[31] S. L. Kukreja, J. Lofberg, and M. J. Brenner, "A least absolute shrinkage and selection operator (lasso) for nonlinear system identification," in *Proceedings of International Federation of Automatic Control Symposium on System Identification*, 2006, pp. 814–819.

[32] M. Sipser and D. Spielman, "Expander codes," *IEEE transactions on Information Theory*, vol. 42, no. 6, pp. 1710–1722, 1996.

[33] W. Xu and B. Hassibi, "Efficient compressive sensing with deterministic guarantees using expander graphs," in *Proceedings of IEEE Information Theory Workshop*, 2007, pp. 414–419.

[34] R. Berinde, A. Gilbert, P. Indyk, H. Karloff, and M. Strauss, "Combining geometry and combinatorics: a unified approach to sparse signal recovery," in *Proceedings of Allerton conference on Communication, Control, and Computing*, 2008, pp. 798–805.